arXiv:1602.05365v1 [cs.PL] 17 Feb 2016

A Specification of Open Transactional Memory for Haskell

Marino Miculan* Marco Peressotti
marino.miculan@uniud.it marco.peressotti@uniud.it
Laboratory of Models and Applications of Distributed Systems
Department of Mathematics, Informatics, and Physics
University of Udine, Italy

Abstract

Transactional memory (TM) has emerged as a promising abstraction for concurrent programming alternative to lock-based synchronizations. However, most TM models admit only *isolated* transactions, which are not adequate in multi-threaded programming where transactions have to interact via shared data *before* committing. In this paper, we present *Open Transactional Memory* (OTM), a programming abstraction supporting *safe, data-driven* interactions between *composable* memory transactions. This is achieved by relaxing isolation between transactions, still ensuring atomicity: threads of different transactions can interact by accessing shared variables, but then their transactions have to commit together—actually, these transactions are transparently *merged*. This model allows for *loosely-coupled* interactions since transaction merging is driven only by accesses to shared data, with no need to specify participants beforehand. In this paper we provide a specification of the OTM in the setting of Concurrent Haskell, showing that it is a conservative extension of current STM abstraction. In particular, we provide a formal semantics, which allows us to prove that OTM satisfies the *opacity* criterion.

1 Introduction

The advent of multicore architectures has emphasized the importance of abstractions supporting correct and scalable multi-threaded programming. In this model, threads can collaborate by interacting on data structures (such as tables, message queues, buffers, etc.) kept in shared memory. Traditional lock-based mechanisms (like semaphores and monitors) used to regulate access to these shared data are notoriously difficult and error-prone, as they easily lead to deadlocks, race conditions and priority inversions; moreover, they are not composable and hinder parallelism, thus reducing efficiency and scalability. *Transactional memory* (TM) has emerged as a promising abstraction to replace locks [3, 14]. The basic idea is to mark blocks of code as *atomic*; then, execution of each block will appear either as if it was executed sequentially and instantaneously at some unique point in time, or, if aborted, as if it did not execute at all. This is obtained by means of *optimistic* executions: the blocks are allowed to run concurrently, and eventually if an interference is detected a transaction is restarted and its effects are rolled back. Thus, each transaction can be viewed in isolation as a *single-threaded* computation, significantly reducing the programmer's burden. Moreover, transactions are composable and ensure absence of deadlocks and priority inversions, automatic roll-back on exceptions, and increased concurrency.

*Partially supported by MIUR project 2010LHT4KM (*CINA*).

However, in multi-threaded programming different transactions may need to interact and exchange data *before* committing. In this situation, transaction isolation is a severe shortcoming. A simple example is a request-response interaction between two transactions via a shared buffer, like in a master/worker situation. We could try to synchronize the threads accessing the buffer `b` by means of two semaphores `c1`, `c2` as follows:

```
// Party1 (Master)
atomically {
  <put request in b>
  up(c1);
  <some other code; may abort>
  down(c2); // wait for answer
  <get answer from b; may abort>
}
```

```
// Party2 (Worker)
atomically {
  down(c1); // wait for data
  <get request from b>
  <compute answer; may abort>
  <put answer in b>
  up(c2);
}
```

Unfortunately, this solution does not work: any admissible execution requires an interleaved scheduling between the two transactions, thus violating isolation; hence, the transactions deadlock as none of them can progress. It is important to notice that this deadlock arises because interaction occurs between threads of *different* transactions; in fact, the solution above is perfectly fine for threads outside transactions or within the same transaction.

To overcome this limitation, in this paper we propose a programming model for *safe, data-driven* interactions between memory transactions. The key observation is that *atomicity* and *isolation* are two disjoint computational aspects:

- an *atomic non-isolated* block is executed "all-or-nothing", but its execution can overlap others' and *uncontrolled* access to shared data is allowed;
- a *non-atomic isolated* block is executed "as it were the only one" (i.e., in mutual exclusion with others), but no rollback on errors is provided.

Thus, a "normal" block of code is neither atomic nor isolated; a mutex block (like Java *synchronized* methods) is isolated but not atomic; and a usual TM transaction is a block which is both atomic and isolated. Our claim is that *atomic non-isolated blocks can be fruitfully used for implementing safe composable interacting memory transactions*—henceforth called *open transactions*.

In this model, a transaction is composed by several threads, called *participants*, which can cooperate on shared data. A transaction commits when all its participants commit, and aborts if any thread aborts. Threads participating to different transactions can access to shared data, but when this happens the transactions are *transparently merged* into a single one. For instance, the two transactions of the synchronization example above would automatically merge becoming the same transaction, so that the two threads can synchronize and proceed. Thus, this model relaxes the isolation requirement still guaranteeing atomicity and consistency; moreover, it allows for *loosely-coupled* interactions since transaction merging is driven only by run-time accesses to shared data, without any explicit coordination among the participants beforehand.

In summary, the contributions of this paper are the following:

- We present *Open Transactional Memory*, a transactional memory model where multi-threaded transactions can interact by non-isolated access to shared data. Consistency and atomicity are ensured by transparently *merging* transactions at runtime.
- We describe this model in the context of Concurrent Haskell (Section 3). Namely, we define two monads `OTM` and `ITM`, representing the computational aspects of atomic *multi-threaded open* (i.e., non-isolated) transactions and atomic *single-threaded isolated* transactions, respectively. Using the construct `atomic`, programs in the `OTM` monad are executed "all-or-nothing" but without isolation; hence these transactions can merge at runtime. When

needed, blocks inside transactions can be executed in isolation by using the construct `isolated`. Both OTM and ITM transactions are *composable*, and we exploit Haskell type system to forbid irreversible effects inside these two monads.[1]

- We provide a formal operational semantics of our system (Section 4). This semantics defines clearly the behaviour also in less intuitive situations, and serves as a reference for implementations. Using this semantics we prove that OTM satisfies the *opacity* correctness criterion for transactions [1].

Some concluding remarks and directions for future work are in Section 5.

2 Concurrency in Haskell

Haskell was born as pure lazy functional language; side effects are handled by means of monads [13]. For instance, I/O actions have type `IO a` and can be combined together by the monadic bind combinator `>>=`. Therefore, the function `putChar :: Char -> IO ()` takes a character and delivers an I/O action that, when performed (even multiple times), prints the given character. Besides external inputs/outputs, values of `IO` include operations with side effects on mutable (typed) cells. A cell holding values of type `a` has type `IORef a` and may be dealt with only via the following operations:

```
newIORef   :: a -> IO (IORef a)
readIORef  :: IORef a -> IO a
writeIORef :: IORef a -> a -> IO ()
```

Concurrent Haskell [12] adds support to threads which independently perform a given I/O action as explained by the type of the thread creation function:

```
forkIO :: IO () -> IO ThreadId
```

The main mechanism for safe thread communication and synchronisation are *MVars*. A value of type `MVar a` is mutable location (as for `IORef a`) that is either empty or full with a value of type `a`. There are two fundamental primitives to interact with MVars:

```
takeMVar :: Mvar a -> IO a
putMvar  :: Mvar a -> a -> IO ()
```

The first empties a full location and blocks otherwise whereas the second fills an empty location and blocks otherwise. Therefore, MVars can be seen as one-place channels and the particular case of `MVar ()` corresponds to binary semaphores.

We refer the reader to [11] for an introduction to concurrency, I/O, exceptions, and cross language interfacing (the "awkward squad" of pure, lazy, functional programming).

STM Haskell [2] builds on Concurrent Haskell adding *transactional actions* and a transactional memory for safe thread communication, called *transactional variables* or *TVars* for short.

Transactional actions have type `STM a` and are concatenated using `STM` monadic "bind" combinator, akin I/O actions. A transactional action remains tentative during its execution and (its effect) is exposed to the rest of the system by

```
atomically :: STM a -> IO a
```

[1] In fact, OTM model can be implemented in any programming language, provided we have some means, either static or dynamic, to forbid irreversible effects inside transactions.

which takes an STM action and delivers an I/O action that, when performed, runs the transaction guaranteeing atomicity and isolation with respect to the rest of the system.

Transactional variables have type `TVar a` where `a` is the type of the value held and, like IOrefs, are manipulated via the interface:

```
newTVar   :: a -> STM (TVar a)
readTVar  :: TVar a -> STM a
writeTVar :: TVar a -> a -> STM ()
```

For instance, the following code uses monadic bind to combine a read and write operation on a transactional variable and define a "transactional update":

```
modifyTVar :: TVar a -> (a -> a) -> STM ()
modifyTVar var f = do
   x <- readTVar var
   writeOTVar var (f x)
```

Then, `atomically (modifyTVar x f)` delivers an I/O action that applies `f` to the value held by `x` and updates `x` accordingly—the two steps being executed as a single atomic isolated operation.

The primitives recalled so far cover memory interaction, but STM allows also for *composable blocking*. In STM Haskell, blocking translates in "this thread has been scheduled too early, i.e., the right conditions are not fulfilled (yet)". The programmer can tell the scheduler about this fact by means of the primitive:

```
retry :: STM a
```

The semantics of `retry` is to abort the transaction and re-run it after at least one of the transactional variables it has read from has been updated—there is no point in blindly restarting a transaction.

Finally, transactions can be composed as alternatives by means of

```
orElse :: STM a -> STM a -> STM a
```

which evaluates its first argument, and if this results is a `retry` the second argument is evaluated discarding any effect of the first.

3 Composable open transactions

In this section we present the key ideas of the paper by gradually introducing the primitives from the OTM library, summarised in Figure 1.

Although the OTM model can be implemented in any language, we consider Haskell because its expressive type system offers a perfect environment for studying the ideas of transactional memory. In [2] this has been used to single out computations which can be executed in transactions, i.e. terms which can perform memory effects, from those which can perform irreversible input/output effects. In this paper we refine further this approach by using the type system to separate *isolated* transactions from those which can interact, and hence merged.

The key point is to separate isolation from atomicity. In fact, isolation is a computational aspect which can be *added* to atomic transactions. From this perspective, we distinguish between isolated atomic actions and (non isolated) atomic actions. The former are values of type `ITM a` and the latter of `OTM a`. Each type of actions can be sequentially composed (by the corresponding monadic binders) preserving atomicity and, for the former, isolation.

```
data ITM a
data OTM a
-- henceforth, t is a placeholder for ITM or OTM --

-- Sequencing, do notation ------------------------
(>>=)  :: t a -> (a -> t b) -> t b
return :: a -> t a

-- Running isolated and atomic computations -------
atomic   :: OTM a -> IO a
isolated :: ITM a -> OTM a
retry    :: ITM a
orElse   :: ITM a -> ITM a -> ITM a

-- Exceptions -------------------------------------
throw :: Exception e => e -> t a
catch :: Exception e => t a -> (e -> t a) -> t a

-- Threading --------------------------------------
fork :: OTM () -> OTM ThreadId

-- Transactional memory ---------------------------
data OTVar a
newOTVar     :: a -> ITM (OTVar a)
readOTVar    :: OTVar a -> ITM a
writeOTVar   :: OTVar a -> a -> ITM ()
```

Figure 1: The base interface of OTM.

The function `isolated` takes an isolated atomic action and delivers an atomic action whose effects are guaranteed to be executed in isolation with respect to other actions. Then, `atomic` takes an atomic action and delivers an I/O action that when performed runs a transaction whose effects are kept tentative until it commits. Tentative effects are shared among all non-isolated transactions. Therefore, any value of type `STM a` can be seen as a value of `ITM a` for the I/O they deliver is the same:

```
atomically  = atomic . isolated
```

Isolation OTM supports composable blocking via the primitive `retry`, under STM slogan "a thread that has to be blocked because it has been scheduled too soon". As for STM, retrying a transactional action actually corresponds to block the threads on some condition. Note that `retry :: OTM a` is not a primitive since it can be defined from that of `ITM` as `isolated retry`.

Checks may be declared as follows:

```
check :: Bool -> ITM ()
check b = if b then return () else retry
```

although similar primitives may be implemented at the runtime level in order to use this information in thread scheduling.

OTM provides a mechanism for safe thread communication by means of transactional variables called *OTVars*, similar to STM's TVars but supporting *open* transactions. These variables are

values of type `OTVar a` where `a` is the type of value held. Creating, reading and writing OTVars is done via the interface shown in Figure 1. All these actions are both atomic and isolated as ensured by their type. Therefore, when it comes to actions of type `ITM a`, OTVars are basically TVars; e.g. `modifyTVar` from STM corresponds to:

```
modifyOTVar :: OTVar a -> (a -> a) -> ITM ()
modifyOTVar var f = do
   x <- readOTVar var
   writeOTVar var (f x)
```

From its type it is immediate to see that the update is both atomic and isolated. In fact, read and write operations are glued together by the `>>=` combinator, preserving both properties.

Likewise, invariants on transactional variables can be easily checked by composing reads and checks as follows:

```
assertOTVar :: OTVar a -> (a -> Bool) -> ITM ()
assertOTVar var p = do
   x <- readOTVar var
   check (p x)
```

Blocking A semaphore is a counter with two fundamental operation: `up` which increments the counter and `down` which decrements the counter if it is not zero and blocks otherwise. Semaphores are implemented using OTM as OTVars holding a counter:

```
type Semaphore = OTVar Int
```

Then, `up` and `down` are two trivial atomic and isolated updates, with the latter being guarded by a pre-condition:

```
up :: Semaphore -> ITM ()
up s = modifyOTvar s (1+)

down :: Semaphore -> ITM ()
down s = do
   assertOTVar s (> 0)
   modifyOTVar s (-1+)
```

Actions can also be composed as alternatives by means of the primitive `orElse`. For instance, the following takes a family of semaphores and delivers an action that decrements one of them, blocking only if none can be decremented:

```
downAny :: [Sempahore] -> ITM ()
downAny (x:xs) = down x `orElse` downAny xs
downAny [] = retry
```

Interaction The interchangeability of OTM and STM ends when isolation is dropped. In fact, OTM offers shared OTVars as a mechanism for safe *transaction interaction*. This means that non-isolated transactional actions see the effects on shared variables of any other non-isolated transactional action, as they are performed concurrently on the same object. This flow of information introduces dependencies between concurrent tentative actions tying together their fate: an action cannot make its effects permanent, if it depends on informations produced by another action which fails to complete. OTM guarantees coherence of transactional actions in presence

of interaction through shared transactional variables. Thus, OTVars enables loosely-coupled interaction right inside atomic actions taking the programming style of STM a step further. For instance, communication, rendezvous, brokering, and in general, multi-party interactions can all be atomic (non-isolated) actions.

In order to substantiate these claims, let us see open transactions in action by implementing a synchronisation scenario as described in Section 1. In this example a master process outsources part of an atomic computation to some thread chosen from a worker pool; data is exchanged via some shared variable, whose access is coordinated by a pair of semaphores. Notably, both the master and the worker can abort the computation at any time, leading the other party to abort as well. This can be achieved straightforwardly using OTM:

```
master c1 c2 = do
   -- put request
   isolated (up c1)
   -- do something else
   isolated (down c2)
   -- get answer
```

```
worker c1 c2 = do
   -- do something
   isolated (down c1)
   -- get request
   -- put answer
   isolated (up c2)
```

Both functions deliver atomic actions in `OTM`, and hence are not isolated. We used semaphores for the sake of exposition but we could synchronize by means of more abstract mechanisms, like barriers, channels or futures, which can be implemented using OTM.

Concurrency Differently from STM, OTM supports parallelism inside non-isolated transactions. We can easily fork new threads without leaving `OTM` but, like any effect of a transactional action, thread creation and execution remain tentative until the whole transaction commits. Forked threads participate to their transaction and impact its life-cycle (e.g. issuing aborts) as any other participant. This means that before committing, all forked threads have to complete their transactional action, i.e. terminate. Therefore, although the whole effect delivered by the transaction has happened concurrently, forked threads never leave a transaction alive.

Because of their transactional nature, threads forked inside a transaction do not have compensations nor continuations (i.e. I/O actions to be executed after an abort or after a commit). Compensations are pointless since aborts revert all effects including thread creation. It is indeed possible to replace the primitive `fork` with one supporting I/O actions as continuations like

```
forkCont :: OTM a -> (a -> IO ()) -> OTM ThreadID
```

In fact, this mechanism can be implemented by means of the primitives already offered OTM: since commits are synchronisation points, the above corresponds to the parent thread forking a thread for each continuation, after the atomic action is successfully completed.

On the other hand, by definition isolated atomic actions have to appear as being executed in a single-threaded setting; hence `ITM`, like `STM`, does not support thread creation.

4 Formal specification of OTM

4.1 Syntax and abstract machine states

We fix an Haskell-like language extended with the OTM primitives of Figure 1. The syntax is summarised in Figure 2 where the meta-variables x and r range over a given countable set of variables Var and of location names Loc, respectively. We assume Haskell typing conventions and denote the set of all well-typed terms by Term.

$$
\begin{array}{lrl}
\text{Value} & V ::= & r \mid \backslash x \texttt{ -> } M \mid \texttt{return}\ M \mid M \texttt{ >>= } N \mid \texttt{throw}\ M \mid \texttt{catch}\ M\ N \mid \texttt{putChar}\ c \mid \\
 & & \texttt{getChar} \mid \texttt{fork}\ M \mid \texttt{atomic}\ M\ N \mid \texttt{isolated}\ M \mid \texttt{retry} \mid M\ \texttt{`orElse`}\ N \mid \\
 & & \texttt{newOTVar}\ M \mid \texttt{readOTVar}\ r \mid \texttt{writeOTVar}\ r\ M \\
\text{Term} & M, N ::= & x \mid V \mid M\,N \mid \ldots
\end{array}
$$

Figure 2: The syntax of values and terms.

$$
\begin{array}{lrll}
\text{Thread} & T_t ::= & (\!|M|\!)_t \mid (\!|M; N|\!)_{t,k} & \\
\text{Thread family} & P ::= & T_{t_1} \parallel \cdots \parallel T_{t_n} & \forall i, j\ t_i \neq t_j \\
\text{Expression} & \mathbb{E} ::= & [-] \mid \mathbb{E} \texttt{ >>= } M & \\
\text{Plain process} & \mathbb{P}_t ::= & (\!|\mathbb{E}|\!)_t \parallel P & t \notin P \\
\text{Transaction} & \mathbb{T}_{t,k} ::= & (\!|\mathbb{E}; M|\!)_{t,k} \parallel P & t \notin P \\
\text{Any process} & \mathbb{A}_t ::= & \mathbb{P}_t \mid \mathbb{T}_{t,k} &
\end{array}
$$

Figure 3: Threads and evaluation contexts.

Terms of this language are evaluated by an abstract state machine whose states are pairs $\langle P; \Sigma \rangle$ formed by:

- a *thread family* (process) $P = T_{t_1} \parallel \cdots \parallel T_{t_n}$,
- a *memory* $\Sigma = \langle \Theta, \Delta, \Psi \rangle$, where $\Theta : \mathsf{Loc} \rightharpoonup \mathsf{Term}$ is the *heap* and $\Delta : \mathsf{Loc} \rightharpoonup \mathsf{Term} \times \mathsf{TrName}$ is the *working memory*; TrName is a set of names used to identify active transactions; Ψ is a forest of threads identifiers.

Threads Threads are the smaller unit of execution the machine scheduler operates on; they execute OTM terms and do not have any private transactional memory. A thread outside transactions is represented by $(\!|M|\!)_t$ where M is the term being evaluated and t is a unique *thread identifier* (Figure 3). A thread inside a transaction k is represented by $(\!|M; N|\!)_{t,k}$ where M is the term being evaluated inside the transaction k and N is the term being evaluated as *continuation* after k commits or aborts.

At any time, all thread identifiers are stored in the auxiliary structure Ψ, which is a forest reflecting how threads are forked: if t' has been forked by t while inside k then t' belongs to k too and occurs in Ψ as a child of t.

We shall present thread families borrowing the parallel operator $\parallel$ from process algebra (Figure 3). The operator is associative, commutative and defined only on threads whose thread identifiers are distinct. The notation is extended to thread families (i.e. processes) with $\mathbf{0}$ denoting the empty family.

Memory The memory Σ is divided in the heap Θ and in the distributed working memory Δ (plus the auxiliary structure Ψ recording thread fork hierarchy). As for traditional closed (ACID) transactions (e.g. [2]), operations inside a transaction are evaluated against Δ and effects are propagated to Θ only on commits. When a thread inside a transaction k accesses a location outside Δ the location is *claimed by transaction* k and remains claimed until k commits, aborts or restarts. Threads in k can interact only with locations claimed by k, but active transactions can be merged to share their claimed locations.

We shall denote the set of all possible states as State, and reference to each projected component of Σ by a subscript, i.e. Σ_Θ for the heap and Σ_Δ for the working memory. When describing updates to the memory Σ, we adopt the convention that Σ' has to be intended equals to Σ except

$$[\textsc{Eval}]\ \frac{M \not\equiv V \quad \mathcal{V}[M] = V}{M \to V}$$

$$[\textsc{BindVal}]\ \frac{}{\texttt{return}\ M\ \texttt{>>=}\ N \to N\ M} \qquad [\textsc{BindEx}]\ \frac{\texttt{e} \in \{\texttt{retry}, \texttt{throw}\ N\}}{\texttt{e}\ \texttt{>>=}\ M \to \texttt{e}}$$

$$[\textsc{CatchVal}]\ \frac{\texttt{r} \in \{\texttt{retry}, \texttt{return}\ N\}}{\texttt{r}\ \texttt{`catch`}\ M \to \texttt{r}} \qquad [\textsc{CatchEx}]\ \frac{}{\texttt{throw}\ M\ \texttt{`catch`}\ N \to N\ M}$$

Figure 4: Term reductions: $M \to N$.

$$[\textsc{InChar}]\ \frac{}{\langle \mathbb{P}_t[\texttt{getChar}]; \Sigma\rangle \xrightarrow{?c} \langle \mathbb{P}_t[\texttt{return}\ c]; \Sigma\rangle}$$

$$[\textsc{OutChar}]\ \frac{}{\langle \mathbb{P}_t[\texttt{putChar}\ c]; \Sigma\rangle \xrightarrow{!c} \langle \mathbb{P}_t[\texttt{return}\ ()]; \Sigma\rangle}$$

$$[\textsc{TermIO}]\ \frac{M \to N}{\langle \mathbb{P}_t[M]; \Sigma\rangle \to \langle \mathbb{P}_t[N]; \Sigma\rangle}$$

$$[\textsc{ForkIO}]\ \frac{t' \notin \mathsf{threads}\mathbb{P}_t[\texttt{fork}\ M]}{\langle \mathbb{P}_t[\texttt{fork}\ M]; \Sigma\rangle \to \langle \mathbb{P}_t[\texttt{return}\ t'] \parallel (\!|M; \texttt{return}|\!)_{t',k}; \Sigma\rangle}$$

Figure 5: `IO` state transitions.

if stated otherwise, i.e. by statements like $\Sigma'_\Theta = \Sigma_\Theta[r \mapsto M]$. Finally, the completely undefined partial function $\varnothing$ denotes the empty heap and working memory.

4.2 Operational semantics

The dynamics of the machine is defined by two transition relations presented in Figures 4 to 7. The first relation $M \to N$ is defined on terms only and models pure computations (Figure 4). Rule (Eval) allows a term M that is not a value to be evaluated by means of an auxiliary (partial) function $\mathcal{V}[M]$ yielding the value V; the other rules define the semantics of the monadic `bind` and exception handling in a standard way. We remark the symmetry between `bind` and `catch` and how `retry` is treated as an exception by (BindEx) and as a result value by (CatchVal).

Relation $\to$ can be thought as accessory to the second relation $\langle P; \Sigma\rangle \xrightarrow{\beta} \langle P'; \Sigma'\rangle$, which describes state transitions. Since several rules can apply to a given state according to different evaluation contexts as per Figure 3, this relation is non-deterministic; this models the fact that the scheduler can choose which thread to execute next among various possibilities. Labels β describe the kind of transition, and are defined as follows:

$$\beta ::= \tau \mid new\langle k\rangle \mid co\langle k\rangle \mid ab\langle k, t, M\rangle \mid \overline{ab}\langle k, t, M\rangle \qquad \text{for } k \in \mathsf{TrName},\ M \in \mathsf{Term}$$

Transitions labelled by τ represent *internal* steps of transitions, i.e., steps which do not need a coordination among transactions: reduction of pure terms, thread creation and memory operations. These transitions are defined by the rules in Figure 6. Reading a location falls into

$$[\textsc{TermT}]\ \frac{M \to N}{\langle \mathbb{T}_{t,k}[M]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[N]; \Sigma\rangle}$$

$$[\textsc{ForkT}]\ \frac{t' \notin \mathsf{threads}(\mathbb{T}_{t,k}[\mathtt{fork}\ M]) \qquad \Sigma'_\Psi = \mathsf{add_child}(t, t', \Sigma_\Psi)}{\langle \mathbb{T}_{t,k}[\mathtt{fork}\ M]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[\mathtt{return}\ t'] \parallel (\!|M; \mathtt{return}|\!)_{t',k}; \Sigma'\rangle}$$

$$[\textsc{NewVar}]\ \frac{r \notin \mathrm{dom}(\Sigma_\Theta) \cup \mathrm{dom}(\Sigma_\Delta) \qquad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M, k)]}{\langle \mathbb{T}_{t,k}[\mathtt{newOTVar}\ M]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[\mathtt{return}\ r]; \Sigma'\rangle}$$

$$[\textsc{Read1}]\ \frac{r \notin \mathrm{dom}(\Sigma_\Delta) \quad \Sigma_\Theta(r) = M \quad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M, k)]}{\langle \mathbb{T}_{t,k}[\mathtt{readOTVar}\ r]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[\mathtt{return}\ M]; \Sigma'\rangle}$$

$$[\textsc{Read2}]\ \frac{\Sigma_\Delta(r) = (M, j) \qquad \Sigma'_\Delta = \Sigma_\Delta[k \mapsto j]}{\langle \mathbb{T}_{t,k}[\mathtt{readOTVar}\ r]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,j}[\mathtt{return}\ M]; \Sigma'\rangle}$$

$$[\textsc{Write1}]\ \frac{r \notin \mathrm{dom}(\Sigma_\Delta) \qquad \Sigma'_\Delta = \Sigma_\Delta[r \mapsto (M, k)]}{\langle \mathbb{T}_{t,k}[\mathtt{writeOTVar}\ r\ M]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[\mathtt{return}\ ()]; \Sigma'\rangle}$$

$$[\textsc{Write2}]\ \frac{\Sigma_\Delta(r) = (N, j) \qquad \Sigma'_\Delta = \Sigma_\Delta[k \mapsto j][r \mapsto (M, j)]}{\langle \mathbb{T}_{t,k}[\mathtt{writeOTVar}\ r\ M]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[\mathtt{return}\ ()][k \mapsto j]; \Sigma'\rangle}$$

$$[\textsc{Or1}]\ \frac{\mathtt{op} \in \{\mathtt{throw}, \mathtt{return}\} \qquad \langle (\!|M; \mathtt{return}|\!)_{t,k}; \Sigma\rangle \xrightarrow{\tau}{}^{*} \langle (\!|\mathtt{op}\ N; \mathtt{return}|\!)_{t,j}; \Sigma'\rangle}{\langle \mathbb{T}_{t,k}[M\ \text{`}\mathtt{orElse}\text{`}\ M']; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}'_{t,j}[\mathtt{op}\ N]; \Sigma'\rangle}$$

$$[\textsc{Or2}]\ \frac{\langle (\!|M; \mathtt{return}|\!)_{t,k}; \Sigma\rangle \xrightarrow{\tau}{}^{*} \langle (\!|\mathtt{retry}; \mathtt{return}|\!)_{t,j}; \Sigma'\rangle}{\langle \mathbb{T}_{t,k}[M\ \text{`}\mathtt{orElse}\text{`}\ M']; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,k}[M']; \Sigma\rangle}$$

$$[\textsc{Isolated}]\ \frac{\mathtt{op} \in \{\mathtt{throw}, \mathtt{return}\} \qquad \langle (\!|M; \mathtt{return}|\!)_{t,k}; \Sigma\rangle \xrightarrow{\tau}{}^{*} \langle (\!|\mathtt{op}\ N; \mathtt{return}|\!)_{t,j}; \Sigma'\rangle}{\langle \mathbb{T}_{t,k}[\mathtt{isolated}\ M]; \Sigma\rangle \xrightarrow{\tau} \langle \mathbb{T}_{t,j}[\mathtt{op}\ N]; \Sigma'\rangle}$$

Figure 6: Transactional state transitions: $\langle P; \Sigma\rangle \xrightarrow{\tau} \langle P'; \Sigma'\rangle$.

two cases: rule (Read1) models the reading of an unclaimed location and its memory effect is to record the claim in Δ, while rule (Read2) models the reading of a claimed location and its effect is to merge the transactions of the current thread with that claiming the location. Writes behave similarly. Rules (Or1) and (Or2) describe the semantics of alternative sub-transactions: if the first one `retry`-es the second is executed discarding any effect of the first. Rule (ForkT) spawns a new thread for the current transaction; a term `fork M` can appear inside `atomic`, thus allowing multi-threaded open transactions, but its use inside `isolated` is prevented by the type system and by the shape of (Isolated) as well.

The remaining labels describe state transitions concerning the life-cycle of transactions: creation, commit, abort, and restart (Figure 7). These operations require a coordination among threads; for instance, an abort from a thread has to be propagated to every thread participating to the same transaction. This is captured in the semantics by labelling the transition with the operation and the name of the transaction involved; this information is used by the derivation rules to force synchronisation of all participants of that transaction. To illustrate this mechanism, we describe the commit of a transaction k, namely $\langle P; \Sigma\rangle \xrightarrow{co\langle k\rangle} \langle P'; \Sigma'\rangle$. First, by means of (MCastGroup) we split P into two subprocesses, one of which contains all threads participating in k (those not in k cannot do a transition whose label contains k). Secondly, using

$$[\textsc{New}]\ \frac{}{\langle (\!|\texttt{atomic}\ M\ \texttt{>>=}\ N|\!)_t;\Sigma\rangle \xrightarrow{new\langle k\rangle} \langle (\!|M;N|\!)_{t,k};\Sigma\rangle}$$

$$[\textsc{Commit}]\ \frac{\Sigma'_\Theta = \mathsf{commit}(k,\Sigma) \qquad \Sigma'_\Delta = \mathsf{cleanup}(k,\Sigma)}{\langle (\!|\texttt{return}\ M;N|\!)_{t,k};\Sigma\rangle \xrightarrow{co\langle k\rangle} \langle (\!|\texttt{return}\ M\ \texttt{>>=}\ N|\!)_t;\Sigma'\rangle}$$

$$[\textsc{Abort1}]\ \frac{\begin{array}{c}\Sigma'_\Theta = \mathsf{leak}(k,\Sigma) \qquad \Sigma'_\Delta = \mathsf{cleanup}(k,\Sigma)\\ \Sigma'_\Psi = \mathsf{remove}(r,\Sigma_\Psi) \qquad r = \mathsf{root}(t,\Sigma_\Psi)\end{array}}{\langle (\!|\texttt{throw}\ M;N|\!)_{t,k};\Sigma\rangle \xrightarrow{ab\langle k,t,M\rangle} \langle (\!|\texttt{throw}\ M\ \texttt{>>=}\ N|\!)_t;\Sigma'\rangle}$$

$$[\textsc{Abort2}]\ \frac{\begin{array}{c}\Sigma'_\Theta = \mathsf{leak}(k,\Sigma) \qquad \Sigma'_\Delta = \mathsf{cleanup}(k,\Sigma)\\ \Sigma'_\Psi = \mathsf{remove}(r,\Sigma_\Psi) \qquad r = \mathsf{root}(t,\Sigma_\Psi) \quad r = \mathsf{root}(t',\Sigma_\Psi)\end{array}}{\langle (\!|M';N|\!)_{t',k};\Sigma\rangle \xrightarrow{\overline{ab}\langle k,t,M\rangle} \langle (\!|\texttt{throw}\ M\ \texttt{>>=}\ N|\!)_{t'};\Sigma'\rangle}$$

$$[\textsc{Abort3}]\ \frac{\begin{array}{c}\Sigma'_\Theta = \mathsf{leak}(k,\Sigma) \qquad \Sigma'_\Delta = \mathsf{cleanup}(k,\Sigma)\\ \Sigma'_\Psi = \mathsf{remove}(r,\Sigma_\Psi) \qquad r = \mathsf{root}(t,\Sigma_\Psi) \qquad r \neq \mathsf{root}(t',\Sigma_\Psi)\end{array}}{\langle (\!|M';N|\!)_{t',k};\Sigma\rangle \xrightarrow{\overline{ab}\langle k,t,M\rangle} \langle (\!|\texttt{retry}|\!)_{t'};\Sigma'\rangle}$$

$$[\textsc{MCastAb}]\ \frac{\langle P;\Sigma\rangle \xrightarrow{ab\langle k,t,M\rangle} \langle P';\Sigma'\rangle \quad \langle Q;\Sigma\rangle \xrightarrow{\overline{ab}\langle k,t,M\rangle} \langle Q';\Sigma'\rangle}{\langle P \parallel Q;\Sigma\rangle \xrightarrow{ab\langle k,t,M\rangle} \langle P' \parallel Q';\Sigma'\rangle}$$

$$[\textsc{MCastCo}]\ \frac{\langle P;\Sigma\rangle \xrightarrow{co\langle k\rangle} \langle P';\Sigma'\rangle \quad \langle Q;\Sigma\rangle \xrightarrow{co\langle k\rangle} \langle Q';\Sigma'\rangle}{\langle P \parallel Q;\Sigma\rangle \xrightarrow{co\langle k\rangle} \langle P' \parallel Q';\Sigma'\rangle}$$

$$[\textsc{MCastGroup}]\ \frac{\langle P;\Sigma\rangle \xrightarrow{\beta} \langle P';\Sigma'\rangle \quad \beta \neq \tau \quad \mathsf{transaction}(\beta) \notin \mathsf{transactions}(Q)}{\langle P \parallel Q;\Sigma\rangle \xrightarrow{\beta} \langle P' \parallel Q;\Sigma'\rangle}$$

Figure 7: Transaction management transitions: $\langle P;\Sigma\rangle \xrightarrow{\beta} \langle P';\Sigma'\rangle$.

recursively (MCastCo) we single out every thread in k. Finally, we apply (Commit) provided that every thread is ready to commit, i.e., it is of the form $(\!|\texttt{return}\ M;N|\!)_{t,k}$.

Aborting a transaction works similarly, but it based on vetoes instead of an unanimous vote. Aborts are triggered by unhandled exceptions raised by some thread, but threads react to this situation in different ways:

- threads in the same tree of the thread rasing the exception have been forked within the transaction; hence, the root thread is aborted and all other threads in the tree are killed because their creation, as for any transactional side-effect, have to be discarded;
- threads in different trees joined the transaction after it was created, due to a merging; hence, these threads just retry their transaction, since aborting would require them to handle exceptions raised by "foreign" threads.

Notice that there are no derivation rules for `retry`, since its meaning is to inform the scheduler that the execution is stuck; hence the machine has to re-execute the transaction from the beginning (or a suitable check-point), following a different execution order, if and when possible.

$$\mathsf{threads}(T_{t_1} \parallel \cdots \parallel T_{t_n}) \triangleq \{t_1, \ldots t_n\}$$

$$\mathsf{transaction}(\beta) \triangleq k \text{ for } \beta \in \{new\langle k\rangle, co\langle k\rangle, ab\langle k, t, M\rangle, \overline{ab}\langle k, t, M\rangle\}$$

$$(\Delta[k \mapsto j])(r) \triangleq \begin{cases} \Delta(r) & \text{if } \Delta(r) = (M, l), l \neq k \\ (M, j) & \text{if } \Delta(r) = (M, k) \end{cases}$$

$$\mathsf{transactions}(P) \triangleq \begin{cases} \mathsf{transactions}(P_1) \cup \mathsf{transactions}(P_2) & \text{if } P = P_1 \parallel P_2 \\ \{k\} & \text{if } P = (\!|M; N|\!)_{t,k} \\ \emptyset & \text{otherwise} \end{cases}$$

$$P[k \mapsto j] \triangleq \begin{cases} P_1[k \mapsto j] \parallel P_2[k \mapsto j] & \text{if } P = P_1 \parallel P_2 \\ (\!|M; N|\!)_{t,j} & \text{if } P = (\!|M; N|\!)_{t,k} \\ P & \text{otherwise} \end{cases}$$

$$\Theta[r \mapsto M](s) \triangleq \begin{cases} M & \text{if } r = s \\ \Theta(s) & \text{otherwise} \end{cases}$$

$$\Delta[r \mapsto (M, k)](s) \triangleq \begin{cases} (M, k) & \text{if } r = s \\ \Delta(s) & \text{otherwise} \end{cases}$$

$$\mathsf{cleanup}(k, \Sigma)(r) \triangleq \begin{cases} \bot & \text{if } \Sigma_\Delta(r) = (M, k) \\ \Sigma_\Delta(r) & \text{otherwise} \end{cases}$$

$$\mathsf{commit}(k, \Sigma)(r) \triangleq \begin{cases} M & \text{if } \Sigma_\Delta(r) = (M, k) \\ \Sigma_\Theta(r) & \text{otherwise} \end{cases}$$

$$\mathsf{leak}(k, \Sigma)(r) \triangleq M \text{ if } \Sigma_\Theta(r) = M \text{ or } \Sigma_\Theta(r) = \bot \text{ and } \Sigma_\Delta(r) = (M, k)$$

Figure 8: Auxiliary functions.

4.3 Opacity

In this section we use the formalisation of OTM to prove that it meets the *opacity* criterion.

The opacity correctness criterion for transactional memory [1] is an extension of the classical *serialisability property* for databases with the additional requirement that even non-committed transactions must access consistent states. Intuitively, this property ensures that: *(a)* effects of any committed transaction appear performed at a single, indivisible point during the transaction lifetime, *(b)* effects of any aborted transaction cannot be seen by any other transaction, and *(c)* transactions always access consistent states of the system.

In order to formally capture these intuitive requirements let us recall some notions from [1]. A *history* is a sequence of `read`, `write`, `commit`, and `abort` operations[2] ordered according to the time at which they were issued (simultaneous events are arbitrarily ordered) and such that no operation can be issued by a transaction that has already performed a `commit` or an `abort`. A transaction k is said to be in a history H if the latter contains at least one operation issued by

[2]The definition given in [1] considers finer-grained events; in particular, `read` and `write` operations are formed by `request`, `execution`, and `response` events. However in *loc. cit.* the authors restrict to histories where `request-execution-response` sequences are not interleaved, hence we can consider the simpler `read`/`write`s events in the first place.

k. Any history H defines a *happens-before* partial order $\prec_H$ where $k \prec_H k'$ iff the transaction k becomes committed or aborted in H before k' issues its first operation. If $\prec_H$ is total then H is called *sequential*. For a history H, let *complete*(H) be the set of histories obtained by adding either a commit or an abort for every live transaction in H.

We are now able to recall Guerraoui and Kapałka's definition[3] of opacity.

Definition 4.1 ([1, Def. 1])**.** *A history H is said to be* opaque *if there is a sequential history S equivalent to some history in the set complete(H) such that $\prec_S \subseteq \prec_H$.*

As shown in [1], opacity corresponds to the absence of mutual dependencies between live transactions, where a dependency is created whenever a transaction reads an information written by another or depends from its outcome.

Definition 4.2 (Opacity graph [1, Sec. 5.4])**.** *For a history H let $\ll$ be a total order on the set T of all transactions in H. An* opacity graph *for H and $\ll$, $OPG(H, \ll)$, is a bi-coloured directed graph on T such that a vertex is* red *if the corresponding transaction is either running or aborted, it is* black *otherwise, and such that there is an edge from k to k' whenever any of the following holds:*

(a) k' happens-before k;
(b) k reads something written by k';
(c) k' reads some location written by k and $k' \ll k$;
(d) k' is neither running nor aborted and there are a location r and a transaction k'' such that $k' \ll k''$, k' writes to r, and k'' reads r from k.

The edge is red if the second case applies otherwise it is black. If all edges from red nodes in $OPG(H, \ll)$ are also red then the graph is said to be well-formed*.*

Let H be a history and let k be a transaction appearing in it. A `read` operation by k is said to be *local* (to k) whenever the previous operation by k on the same location was a `write`. A `write` operation by k is said to be *local* (to k) whenever the next operation by k on the same location is a `write`. We denote by *nonlocal*(H) the longest sub-history of H without any local operations. A history H is said *locally-consistent* if every local `read` is preceded by a `write` operation that writes the red value; it is said *consistent* if, additionally, whenever some k reads v from r in *nonlocal*(H) then some k' writes v to r in *nonlocal*(H).

Theorem 4.3 ([1, Thm. 2])**.** *A history H is opaque if and only if (a) H is consistent and (b) there exists a total order $\ll$ on the set of transactions in H such that $OPG(nonlocal(H), \ll)$ is well–formed and acyclic.*

In [1] transactions may encapsulate several threads but cannot be merged. Therefore, in order to study opacity of OTM we extend the set of operations considered in loc. cit. with explicit merges. Let k, k' be two running transactions in the given history; when they merge, they share their threads, locations, and effects. From this perspective, k is commit-pending and depends from k' and hence in the opacity graph, k is a red node connected to k' by a red edge. Hence, merges can be equivalently expressed at the history level by sequences like:

(1) new x; *(2)* k' writes on x; *(3)* k reads from x; *(4)* k prepares to commit.

These are the only dependencies found in histories generated by OTM.

[3]The original definition requires the history H to be also "legal", but this notion is relevant only in presence of non-transactional operations which both STM and OTM prevent by design.

Theorem 4.4. *For H a history describing an execution of a* OTM *program and a total order $\ll$, $OPG(nonlocal(H), \ll)$ is a forest of red edges where only roots may be white.*

Proof. By inspection of the rules it is easy to see that *(a)* transactions may access only locations they claimed; *(b)* claimed locations are released only on `commit`s, `abort`s and retries; *(c)* transactions have to merge with any transaction holding a location they need. Therefore, at any time there is at most one running transaction issuing operations on a given location, hence `read`s and `write`s do not create edges. Thus edges are created only during the execution of merges and, by inspecting the above implementation, it easy to see that *(d)* any transaction can issue at most one merge; *(e)* a transaction issuing a merge is a red node; *(f)* the edge created by a merge is red. Therefore, transactions form a forest made of red edges where any non-root node is red. □

Corollary 4.5 (Opacity)**.** OTM *meets the opacity criterion.*

Proof. A forest formed by red edges whose sources are always red is acyclic and well-formed. □

5 Conclusions

In this paper we have presented OTM, a programming model supporting interactions between composable memory transactions. This model separates isolated transactions from non-isolated ones, still guaranteeing atomicity; the latter can interact by accessing to shared variables. Consistency is ensured by transparently *merging* interacting transactions at runtime. We have showed the versatility and simplicity of OTM by implementing some examples which are incompatible with isolation, and we have given a formal semantics for OTM, which allowed us to prove that this model satisfies the opacity criterion.

There are two main directions for future work each posing its own challenges. First, like STM, this model supports nesting (via `orElse`); however, this feature is currently limited to isolated (sub)transactions. Supporting nesting of open transaction requires additional care in the handling of side-effects: is merging transactions at different level of nesting feasible and meaningful or are we breaking the intuition behind the programming model? Secondly, an implementation is due in order to validate experimentally the model. A possible approach is to implement OTM completely in Haskell on top of STM. This solution does not need any specific support from the Haskell RunTime (HRT) but cannot benefit of the performance gains offered by a deeper integration, thus hindering any fair comparison with existing TM models, like STM. On the other hand, integrating OTM with the HRT and the Glasgow Haskell Compiler, akin STM, would be more efficient but also more complex and invasive.

We have presented OTM within Haskell (especially to leverage its type system), but this model is general and can be applied to other languages. A possible future work is to port this model to an imperative object oriented language, such as Java or C++; however, like other TM implementations, we expect that this extension will require some changes in the compiler and/or the runtime.

This work builds on the ideas in [10] where we described an abstract calculus with shared memory and open transactions. In *loc. cit.* we showed how this model is expressive enough to represent $TCCS^m$ [4], a variant of the Calculus of Communicating Systems with transactional synchronization. Being based on CCS, communication in $TCCS^m$ is synchronous; however, nowadays asynchronous models play an important rôle (see e.g. actors, event-driven programming, etc.). It may be interesting to generalize the discussion so as to consider also this case, e.g. by defining an actor-based calculus with open transactions. Such a calculus can be quite useful also for modelling speculative reasoning for cooperating systems [5–9]. A local version of actor-based open transactions can be implemented in OTM using lock-free data structures (e.g., message queues) in shared transactional memory.

Acknowledgements We thank Nicola Gigante and Valentino Picotti for their valuable feedback about the OTM programming model.

References

[1] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *Proc. PPoPP*, pages 175–184. ACM, 2008.

[2] T. Harris, S. Marlow, S. L. Peyton Jones, and M. Herlihy. Composable memory transactions. In *Proc. PPoPP*, pages 48–60. ACM, 2005.

[3] M. Herlihy and J. E. B. Moss. Transactional memory: Architectural support for lock-free data structures. In A. J. Smith, editor, *Proc. ISCA*, pages 289–300. ACM, 1993.

[4] V. Koutavas, C. Spaccasassi, and M. Hennessy. Bisimulations for communicating transactions. In A. Muscholl, editor, *Proc. FOSSACS*, volume 8412 of *LNCS*, pages 320–334. Springer, 2014.

[5] J. Ma, K. Broda, R. Goebel, H. Hosobe, A. Russo, and K. Satoh. Speculative abductive reasoning for hierarchical agent systems. In J. Dix, J. Leite, G. Governatori, and W. Jamroga, editors, *Computational Logic in Multi-Agent Systems*, volume 6245 of *LNCS*, pages 49–64. 2010.

[6] A. Mansutti, M. Miculan, and M. Peressotti. Distributed execution of bigraphical reactive systems. *ECEASST*, 71, 2014.

[7] A. Mansutti, M. Miculan, and M. Peressotti. Towards distributed bigraphical reactive systems. In R. Echahed, A. Habel, and M. Mosbah, editors, *Proc. GCM'14*, page 45, 2014.

[8] A. Mansutti, M. Miculan, and M. Peressotti. Multi-agent systems design and prototyping with bigraphical reactive systems. In K. Magoutis and P. Pietzuch, editors, *Proc. DAIS*, volume 8460 of *LNCS*, pages 201–208. 2014.

[9] M. Miculan and M. Peressotti. A CSP implementation of the bigraph embedding problem. *CoRR*, abs/1412.1042, 2014.

[10] M. Miculan, M. Peressotti, and A. Toneguzzo. Open transactions on shared memory. In *COORDINATION*, volume 9037 of *LNCS*, pages 213–229. Springer, 2015.

[11] S. L. Peyton Jones. Tackling the awkward squad: monadic input/output, concurrency, exceptions, and foreign-language calls in Haskell. *Engineering theories of software construction*, 180:47, 2001.

[12] S. L. Peyton Jones, A. D. Gordon, and S. Finne. Concurrent Haskell. In H. Boehm and G. L. S. Jr., editors, *Proc. POPL*, pages 295–308. ACM, 1996.

[13] S. L. Peyton Jones and P. Wadler. Imperative functional programming. In M. S. V. Deusen and B. Lang, editors, *Proc. POPL*, pages 71–84. ACM 1993.

[14] N. Shavit and D. Touitou. Software transactional memory. *Distributed Computing*, 10(2):99–116, 1997.